

Computation Institute

Big process for big data

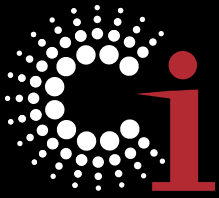
Ian Foster

foster@anl.gov

NASA Goddard, February 27, 2013



computationinstitute.org

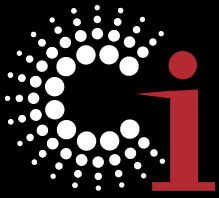


The Computation Institute

= UChicago + Argonne

= Cross-disciplinary nexus

= Home of the Research Cloud



**High energy
physics**

Genetics

Molecular biology

Climate change

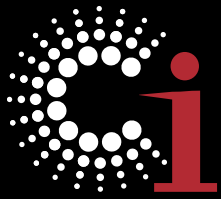
Cosmology

Linguistics

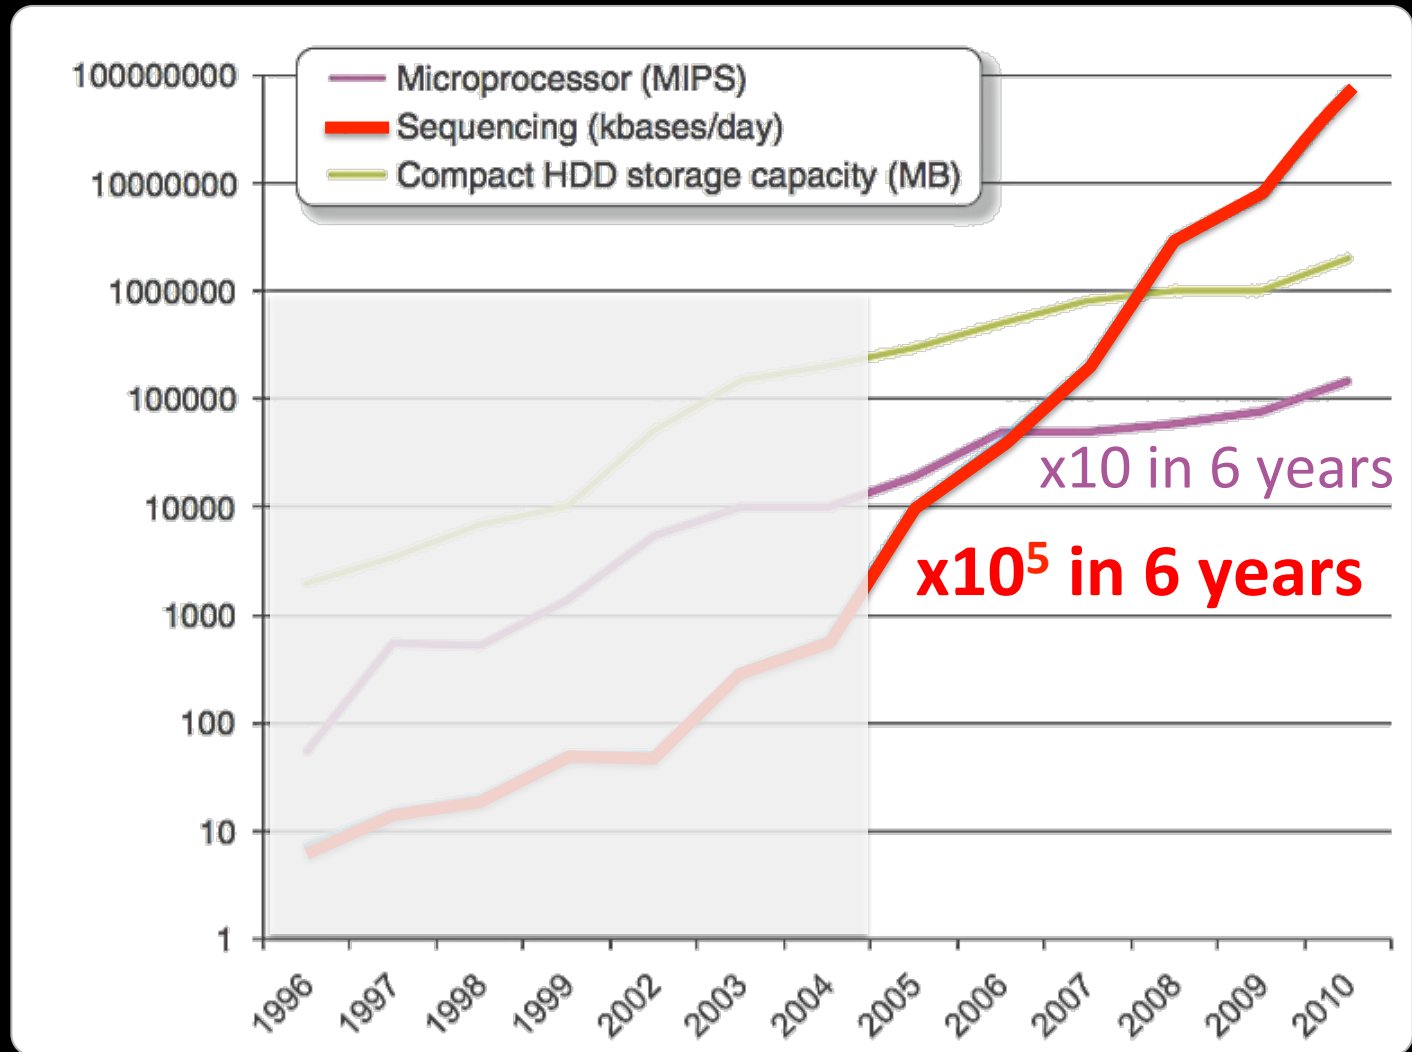
Metagenomics

Economics

Visual arts

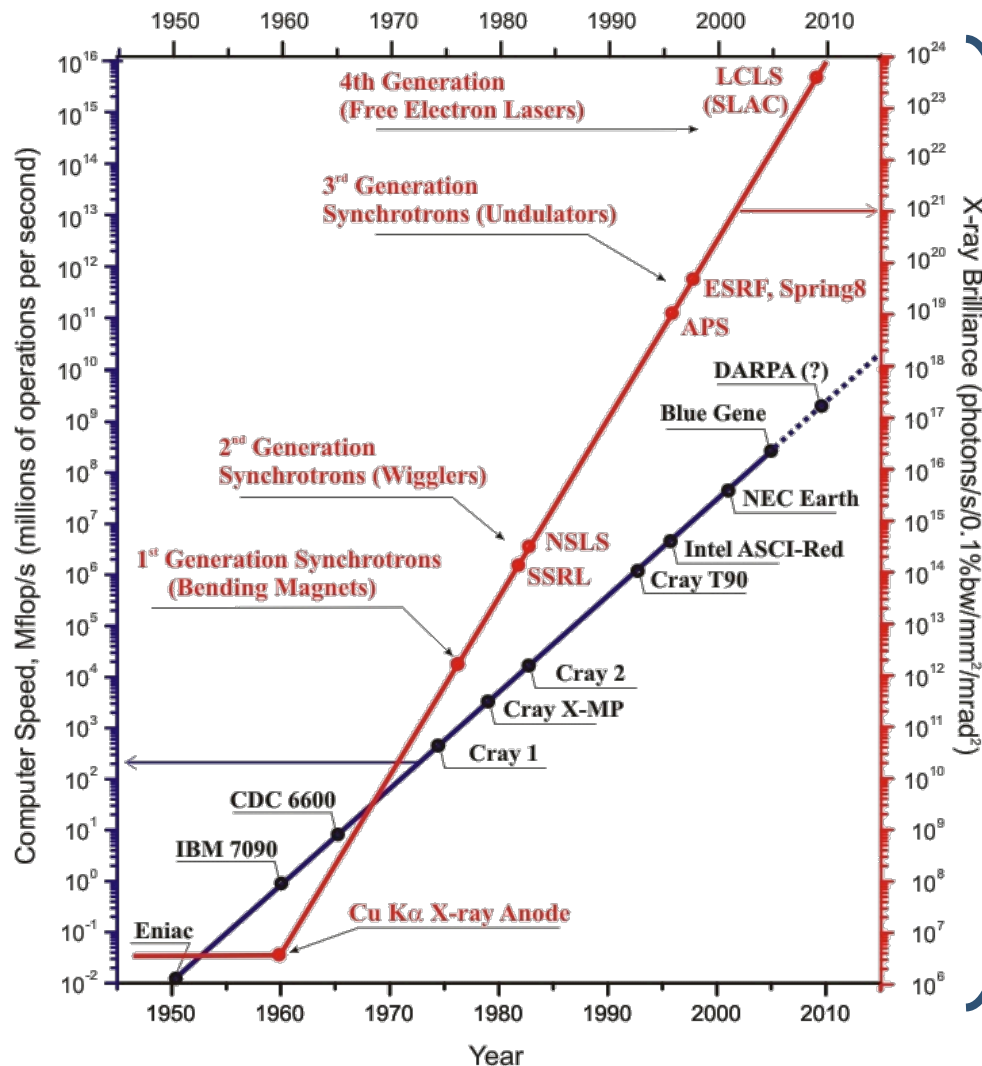


Will data kill genomics?



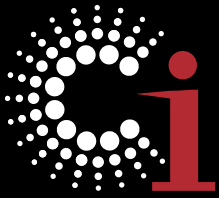


Moore's Law for X-Ray Sources

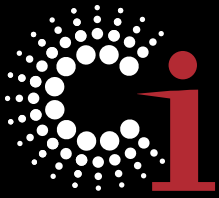


12 orders
of magnitude
in 6 decades!

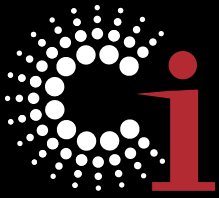
18 orders
of magnitude
in 5 decades!



1.2 PB of climate data
Delivered to 23,000 users

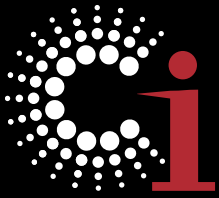


We have exceptional
infrastructure for the 1%



We have exceptional
infrastructure for the 1%

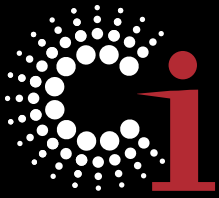
What about the 99%?



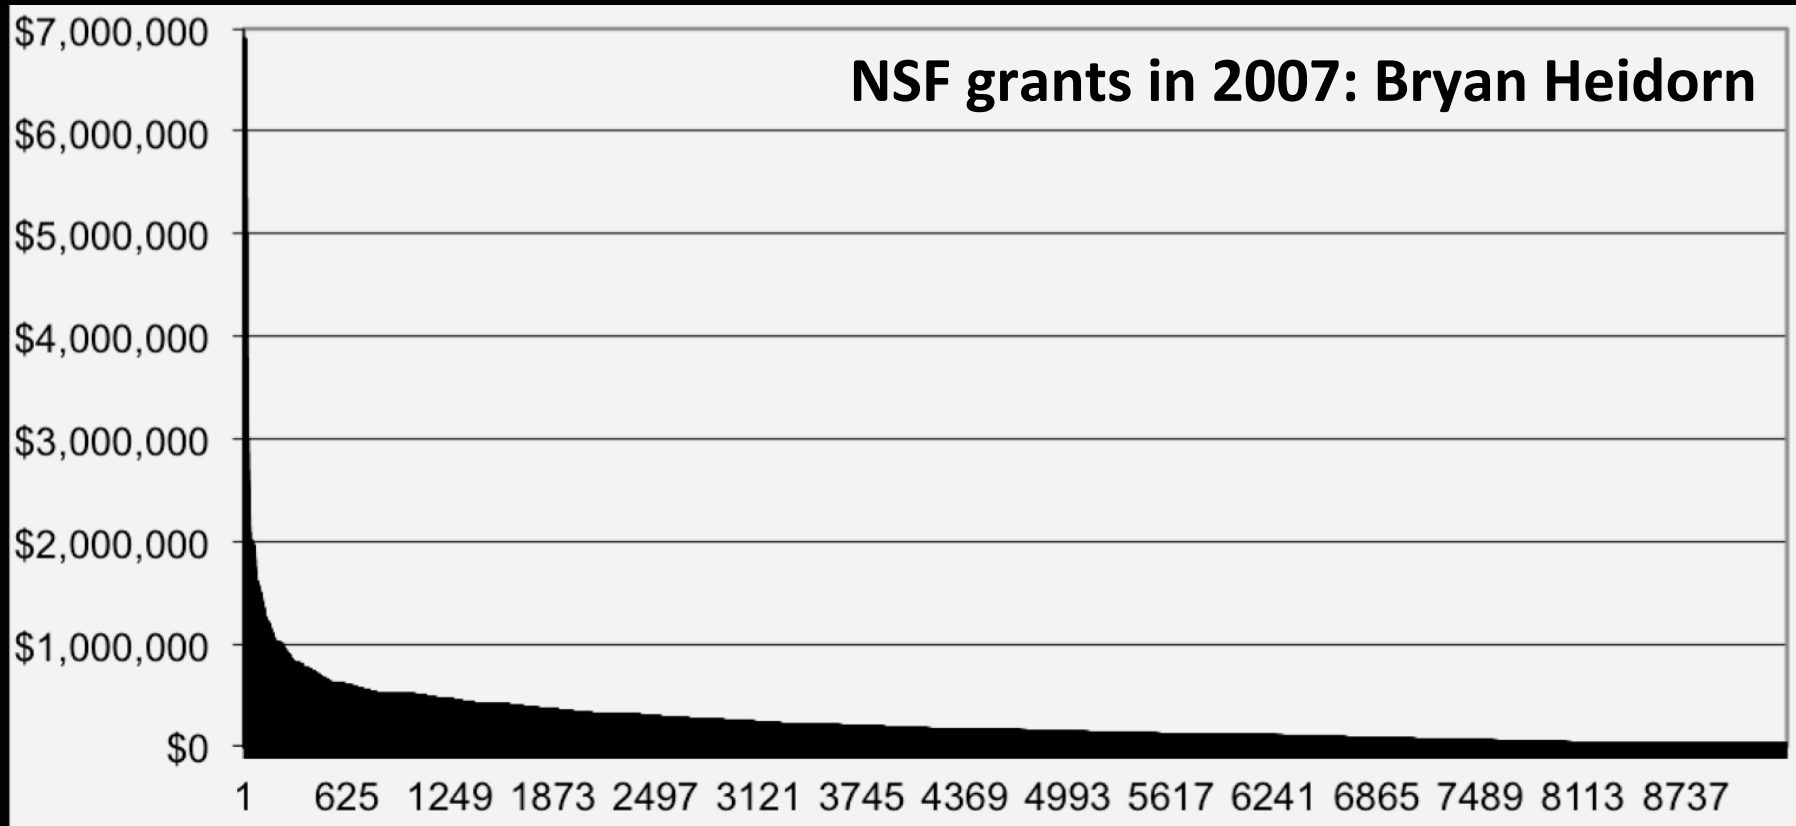
We have exceptional
infrastructure for the 1%

What about the 99%?

Big science. Small labs.



The long tail of research



80% of awards and 50% of \$\$ are for grants < \$350,000

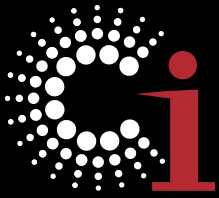


Can they remain competitive?

"Well, in our country," said Alice ...
"you'd generally get to somewhere
else — if you run very fast for a
long time, as we've been doing."

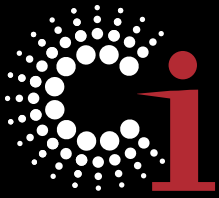
"A slow sort of country!" said the
Queen. "Now, here, you see, it
takes all the running you can do,
to keep in the same place. If you
want to get somewhere else, you
must **run at least twice as fast as that!**"





Need: A new way to deliver
research cyberinfrastructure

Frictionless
Affordable
Sustainable



We asked ourselves:

What if the research work flow
could be managed as easily as...

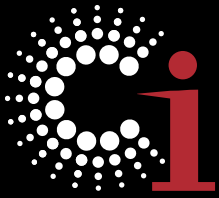
flickr

...our pictures

...our e-mail

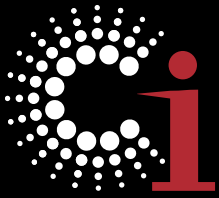


...home entertainment



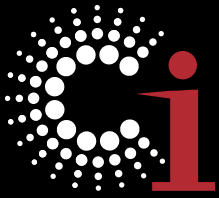
What makes these services great?

Great User Experience
+
High performance
(but invisible) infrastructure



We aspire (initially) to create a
great user experience for
research data management

What would a “dropbox for
science” look like?

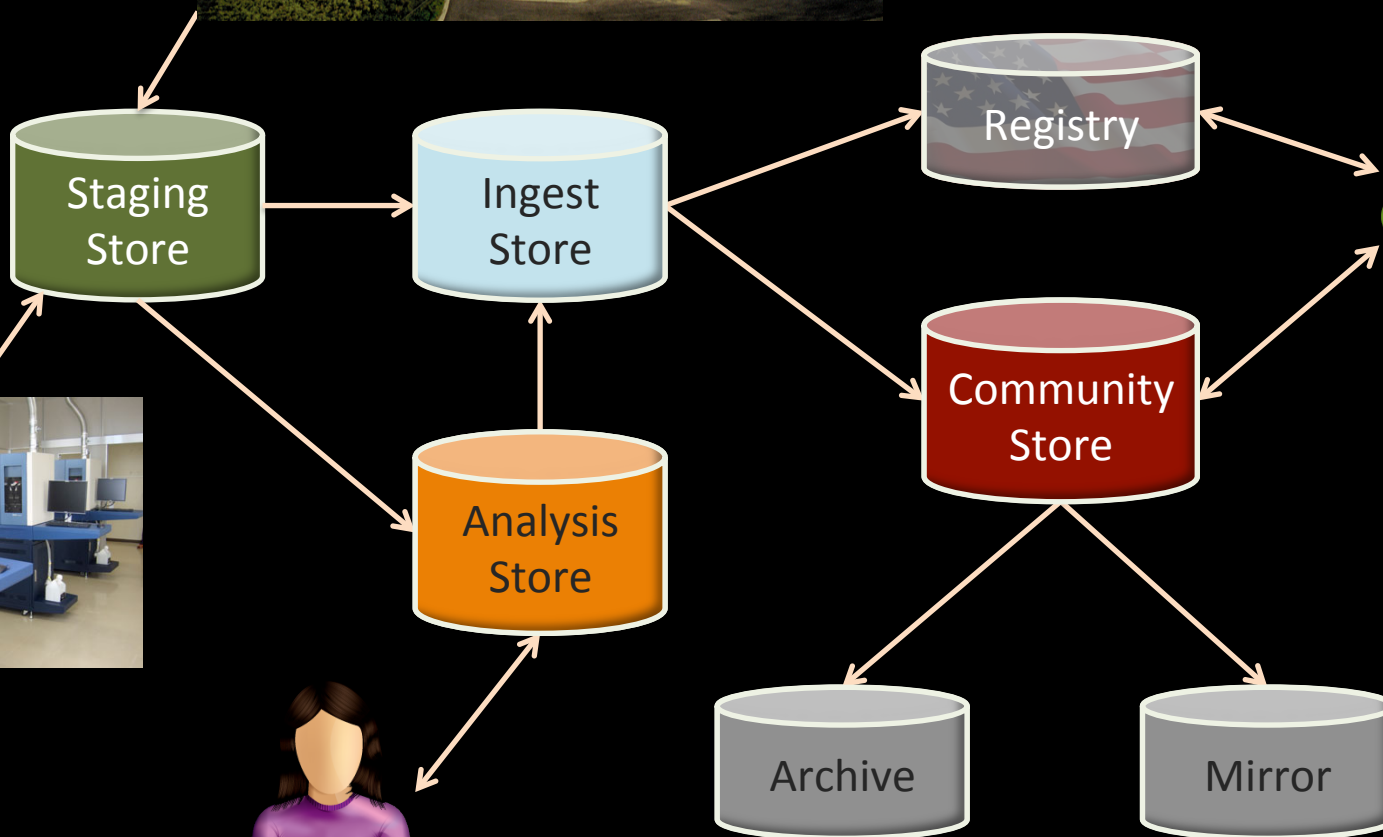


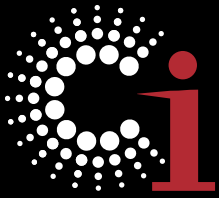
- Collect
- Move
- Sync
- Share
- Analyze
- Annotate
- Publish
- Search
- Backup
- Archive

...for **BIG DATA**



A common work flow...



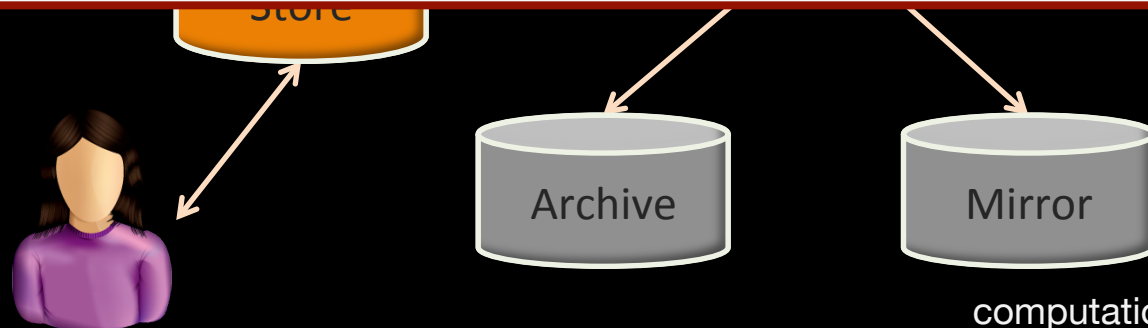


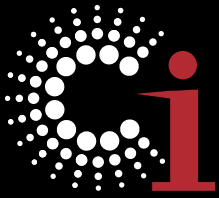
... with common challenges



Data movement, sync, and sharing

- Between facilities, archives, researchers
- Many files, large data volumes
- With security, reliability, performance





- Collect

- Annotate

- **Move**

- **Sync**

- **Share**

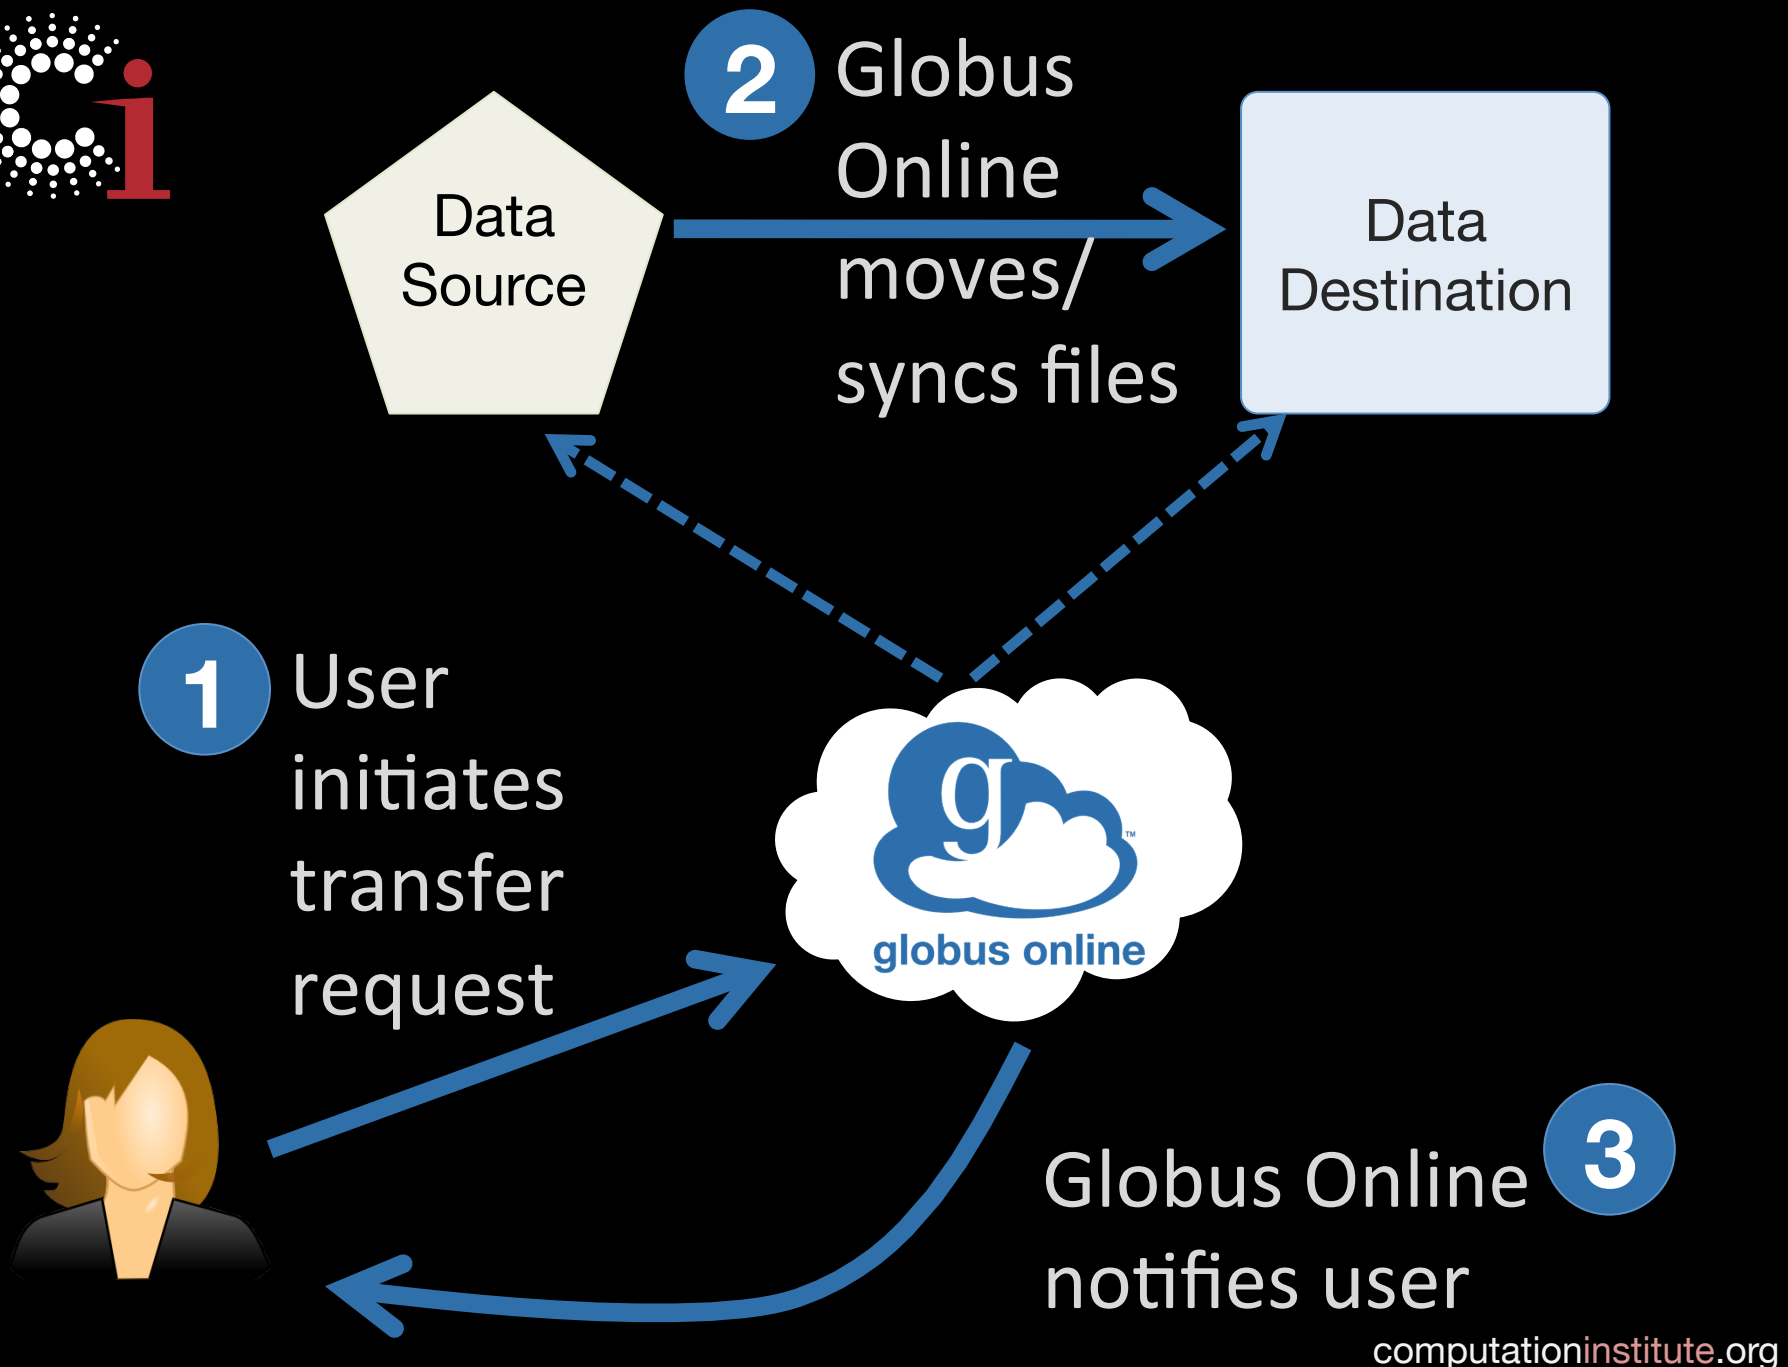
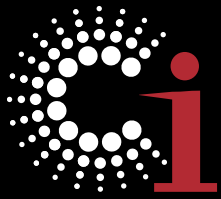


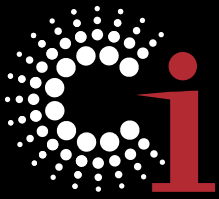
globus online

- Analyze

- Archive

Capabilities delivered using
Software-as-Service (SaaS) model

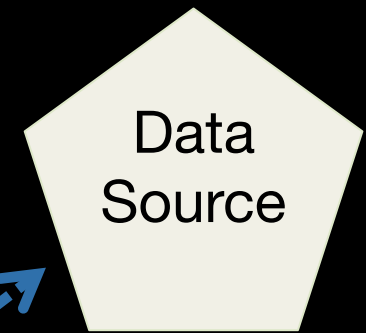




1 User A selects file(s) to share; selects user/group, sets share permissions

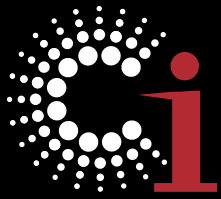


2 Globus Online tracks shared files; no need to move files to cloud storage!



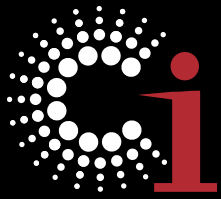
3 User B logs in to Globus Online and accesses shared file





Extreme ease of use

- InCommon, Oauth, OpenID, X.509, ...
- Credential management
- Group definition and management
- Transfer management and optimization
- Reliability via transfer retries
- Web interface, REST API, command line
- One-click “Globus Connect” install
- 5-minute Globus Connect Multi User install

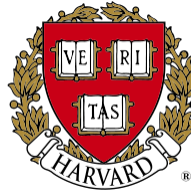


Early adoption is encouraging

XSEDE

Extreme Science and Engineering
Discovery Environment

NERSC



UNIVERSITY OF
EXETER

W
UNIVERSITY of
WASHINGTON



**Carnegie
Mellon
University**



APS
physics



 **THE UNIVERSITY
OF AUCKLAND**
NEW ZEALAND

Te Whare Wānanga o Tāmaki Makaurau



Information Sciences Institute

KSU



INDIANA
UNIVERSITY

 **Fermilab**



EMORY



 **Los Alamos**
NATIONAL LABORATORY
EST. 1943

CORNELL
UNIVERSITY

Cal

Ole Miss



THE UNIVERSITY OF
CHICAGO



NEW YORK UNIVERSITY

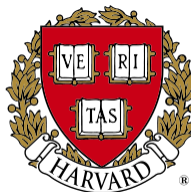
Argonne
NATIONAL LABORATORY



Early adoption is encouraging

XSEDE

Extreme Science and Engineering
Discovery Environment



8,000 registered users; ~100 daily

~10 PB moved; ~1B files

10x (or better) performance vs. scp

99.9% availability

Entirely hosted on AWS

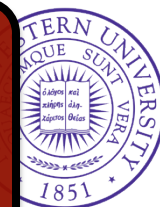
UNIVERSITY OF
WASHINGTON



NEW ZEALAND
Te Whare Wānanga o Te Aroaro

INDIANA

UNIVERSITY



NATIONAL LABORATORY
EST. 1943

CORNELL
UNIVERSITY



Ole Miss

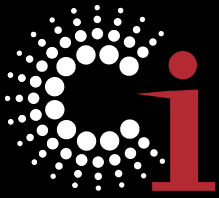


THE UNIVERSITY OF
CHICAGO

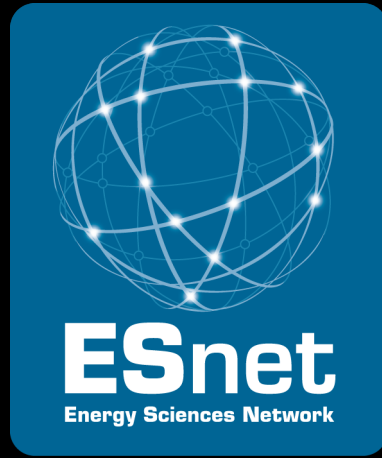
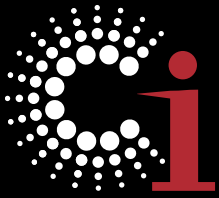


NEW YORK UNIVERSITY





Delivering a great user
experience relies on
**high performance network
infrastructure**



Science DMZ + optimizes performance

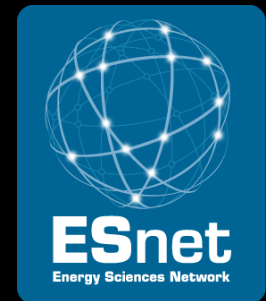
Bandwidth Requirements to move Y Bytes of data in Time X

Bits per Second Requirements

	1H	8H	24H	7Days	30Days
10PB	25,020.0 Gbps	3,127.5 Gbps	1,042.5 Gbps	148.9 Gbps	34.7 Gbps
1PB	2,502.0 Gbps	312.7 Gbps	104.2 Gbps	14.9 Gbps	3.5 Gbps
100TB	244.3 Gbps	30.5 Gbps	10.2 Gbps	1.5 Gbps	339.4 Mbps
10TB	24.4 Gbps	3.1 Gbps	1.0 Gbps	145.4 Mbps	33.9 Mbps
1TB	2.4 Gbps	305.4 Mbps	101.8 Mbps	14.5 Mbps	3.4 Mbps
100GB	238.6 Mbps	29.8 Mbps	9.9 Mbps	1.4 Mbps	331.4 Kbps
10GB	23.9 Mbps	3.0 Mbps	994.2 Kbps	142.0 Kbps	33.1 Kbps
1GB	2.4 Mbps	298.3 Kbps	99.4 Kbps	14.2 Kbps	3.3 Kbps
100MB	233.0 Kbps	29.1 Kbps	9.7 Kbps	1.4 Kbps	0.3 Kbps



What is a Science DMZ?



Three key components, all required:

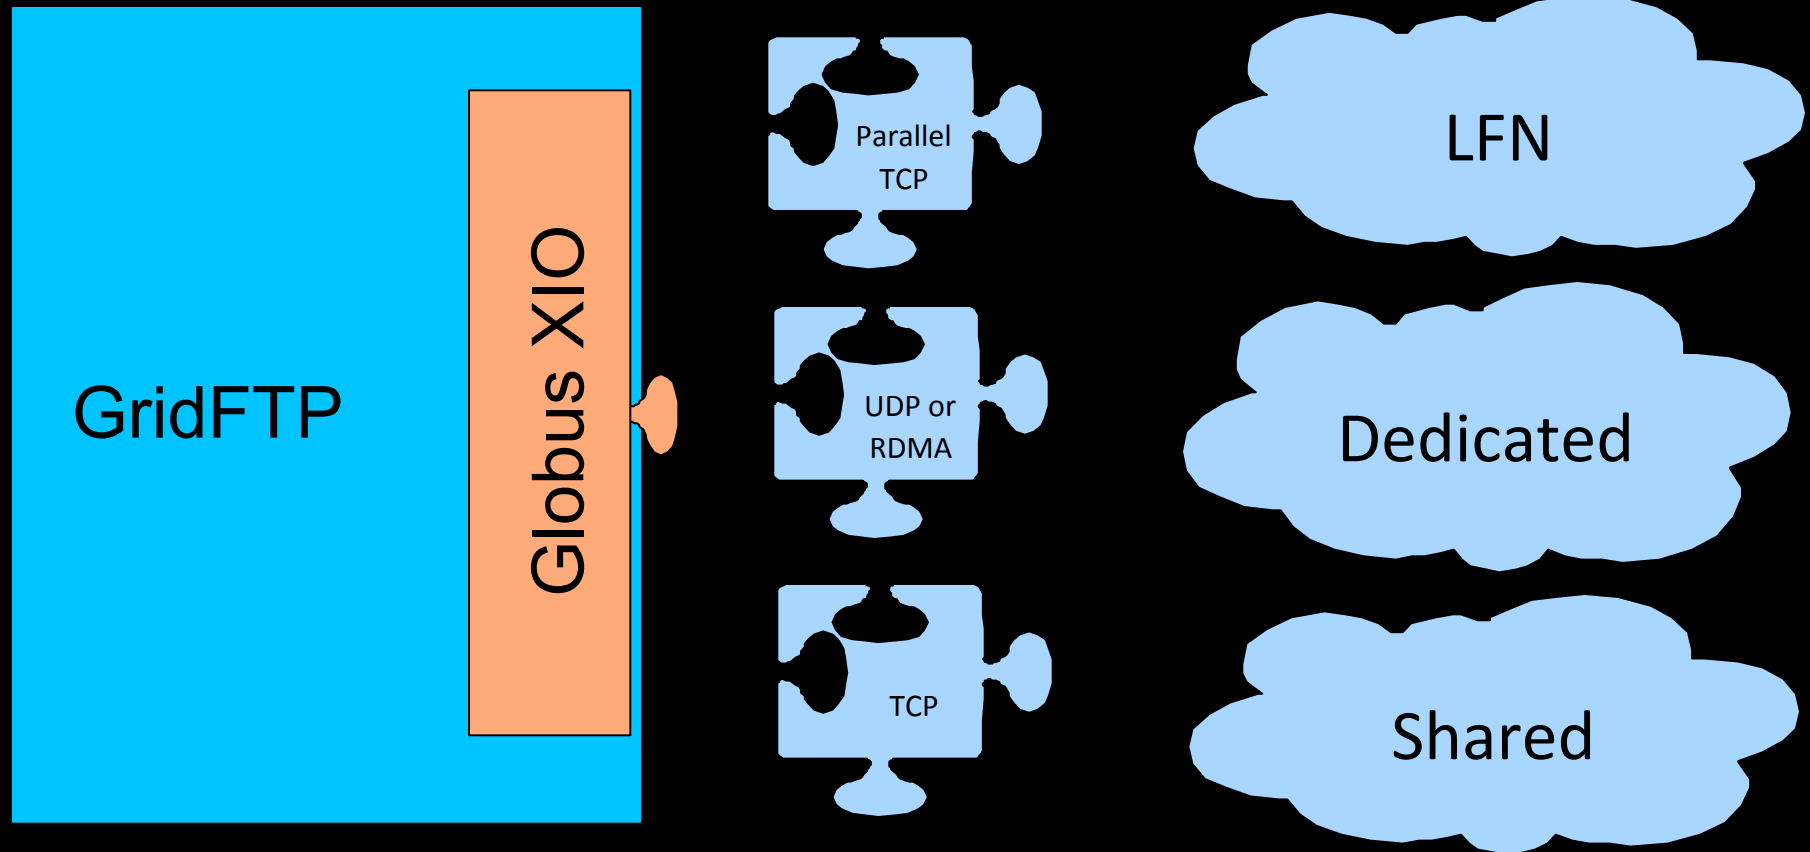
- “Friction free” network path
 - Highly capable network devices (wire-speed, deep queues)
 - Virtual circuit connectivity option
 - Security policy and enforcement specific to science workflows
 - Located at or near site perimeter if possible
- Dedicated, high-performance Data Transfer Nodes (DTNs)
 - Hardware, operating system, libraries optimized for transfer
 - Optimized data transfer tools: Globus Online, GridFTP
- Performance measurement/test node
 - perfSONAR



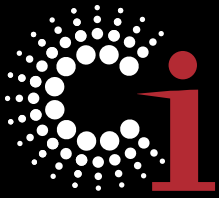
Details at <http://fasterdata.es.net/science-dmz/>



Globus GridFTP architecture



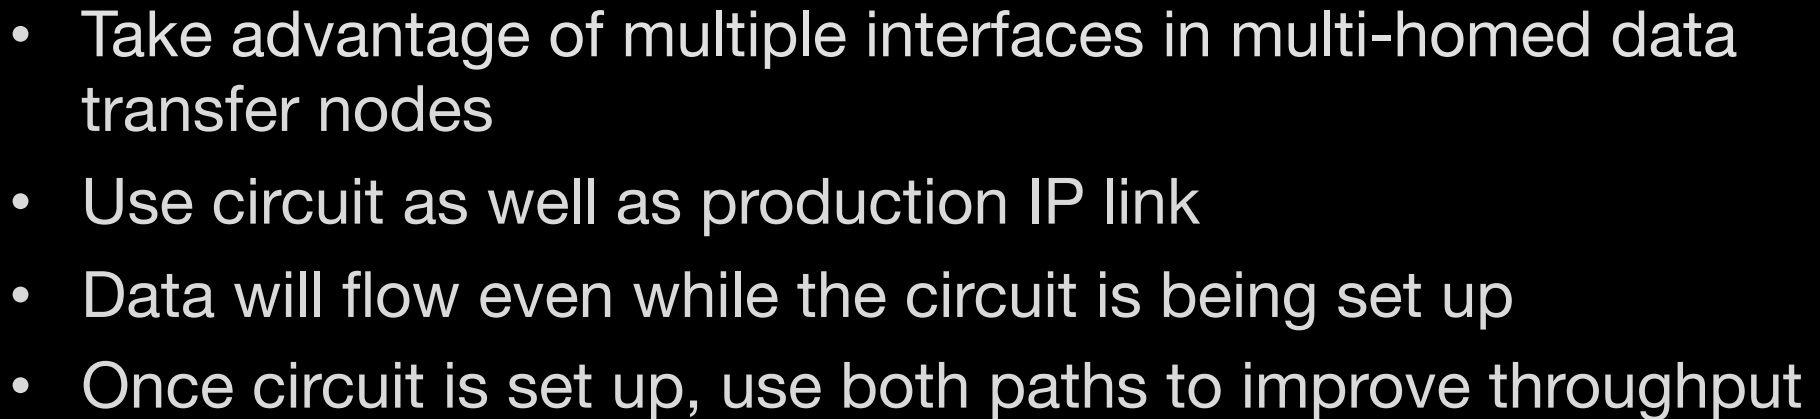
Internal layered XIO architecture allows alternative network and filesystem interfaces to be **plugged in** to the stack

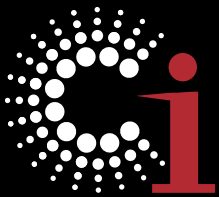


GridFTP performance options

- TCP configuration
- Concurrency: Multiple flows per node
- Parallelism: Multiple nodes
- Pipelining of requests to support small files
- Multiple cores for integrity, encryption
- Alternative protocol selection*
- Use of circuits and multiple paths*

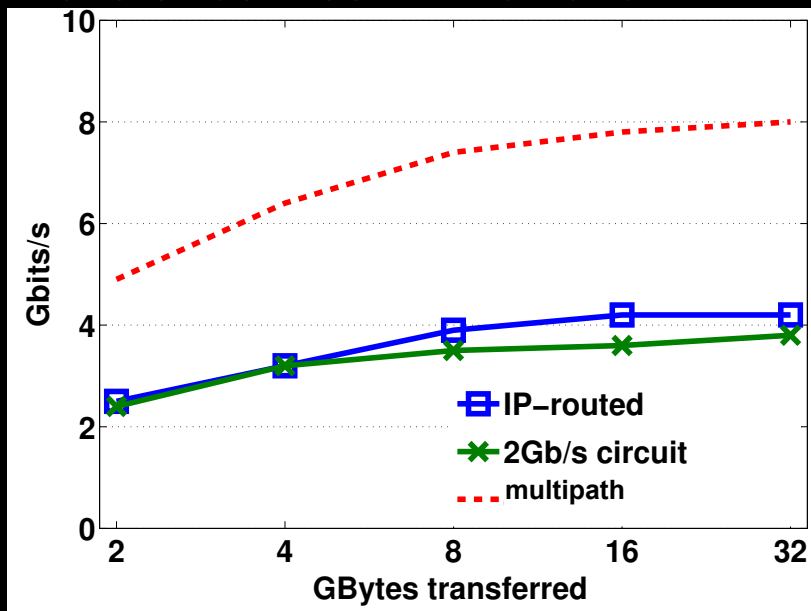
Globus Online can configure these options based on what it knows about a transfer



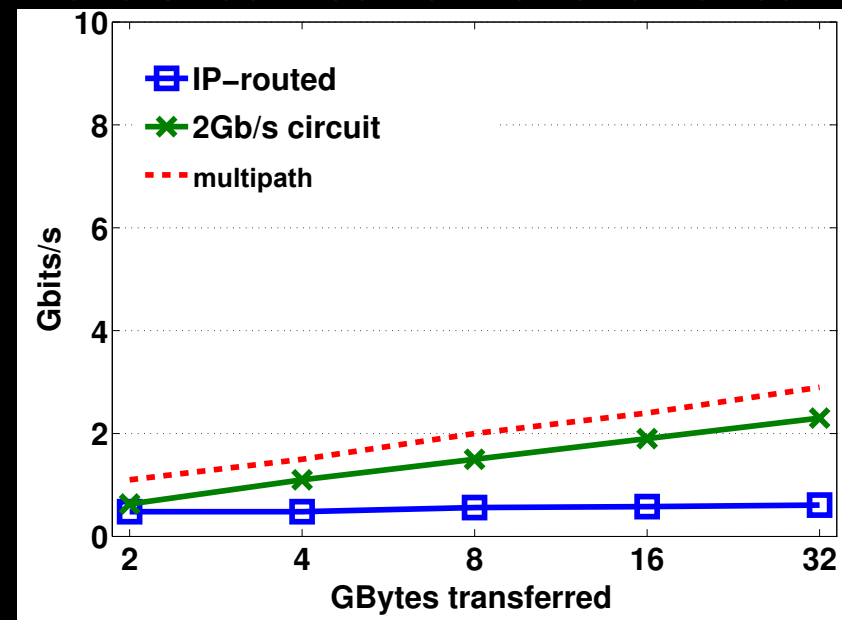


Exploiting multiple paths

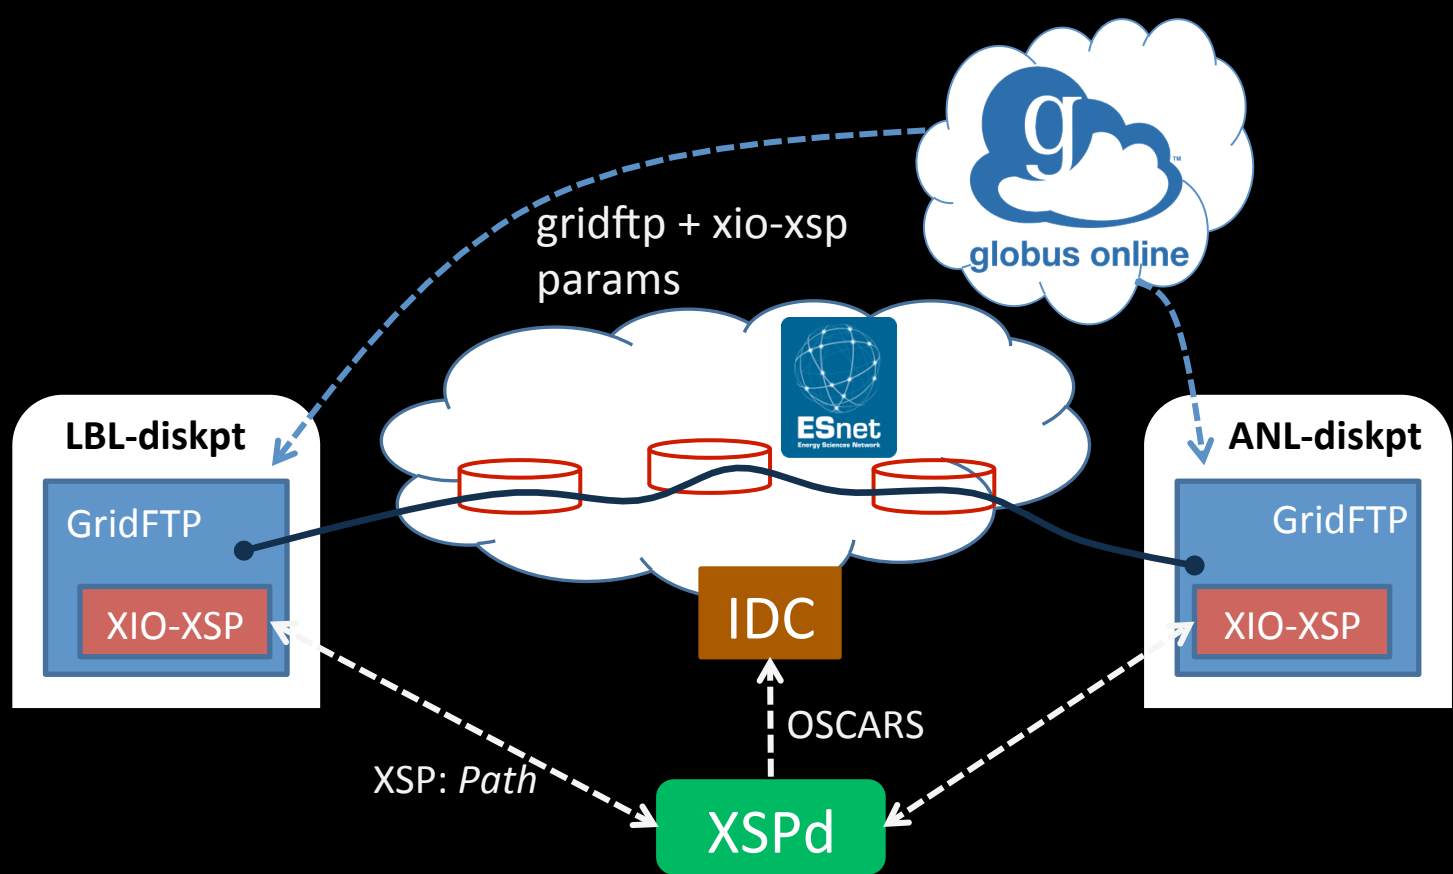
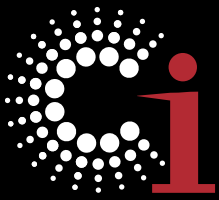
Transfer between NERSC and ANL



Transfer between UMich and Caltech

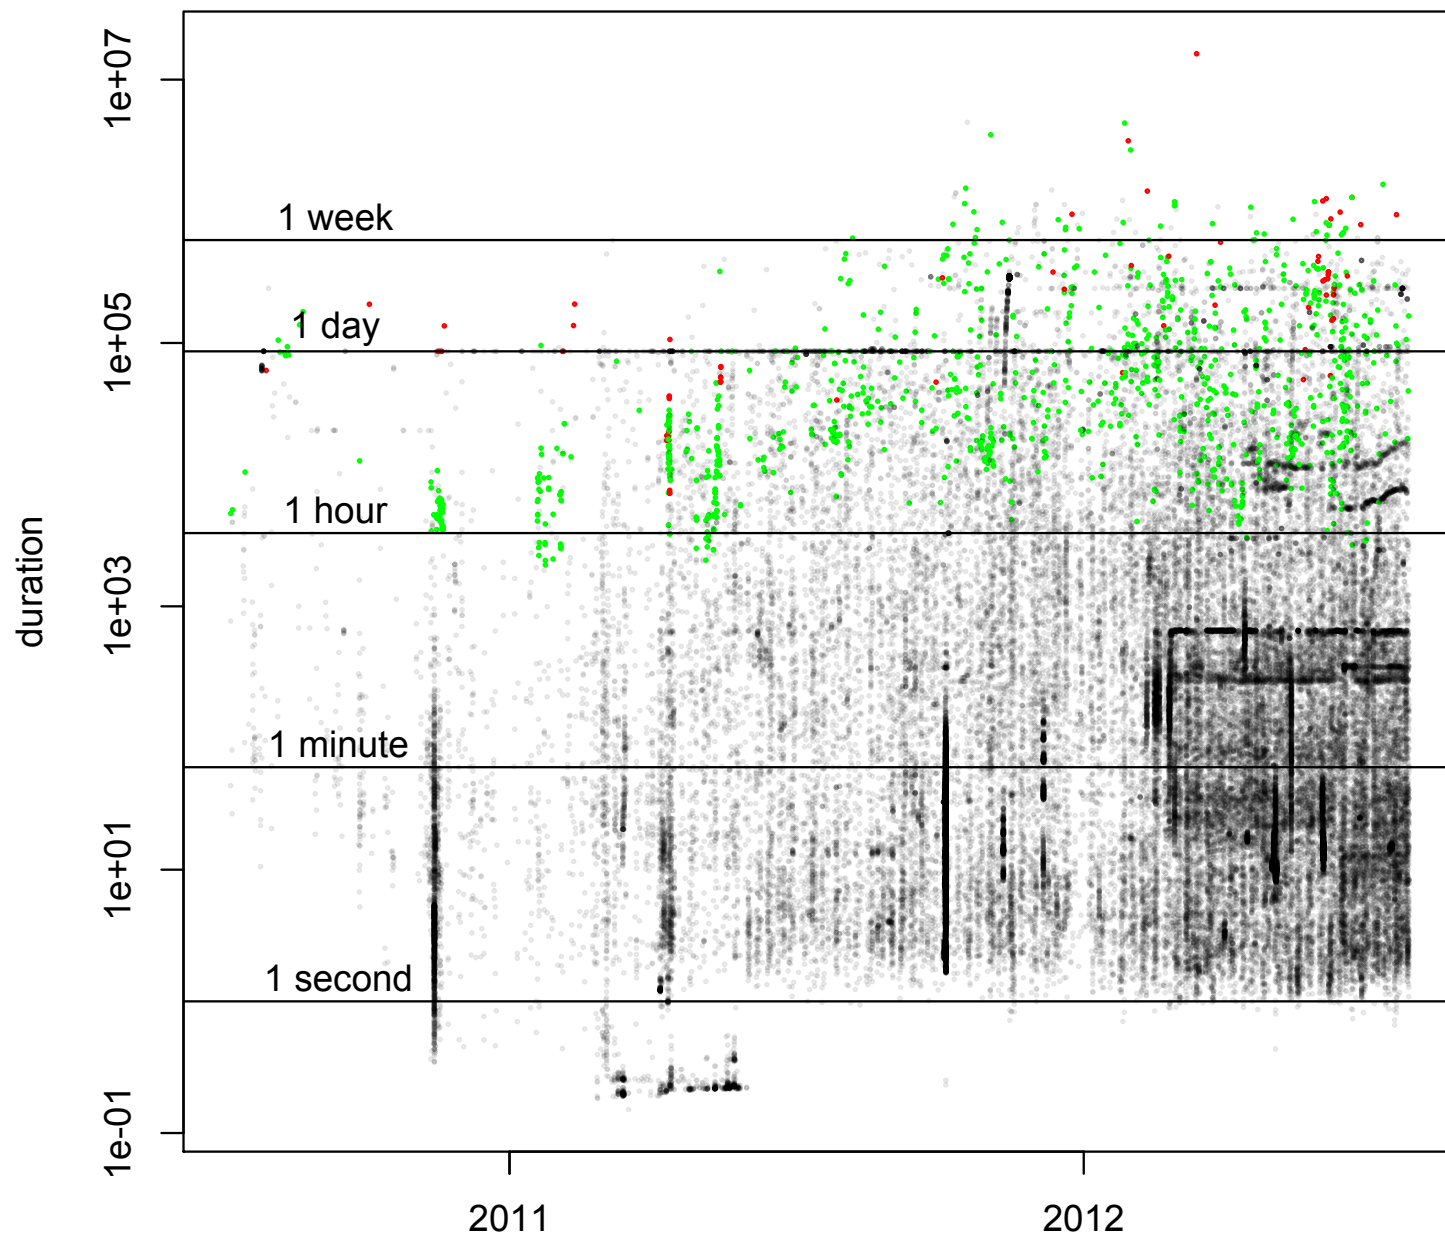


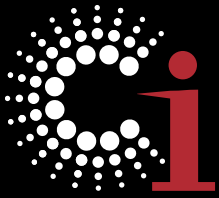
Default, commodity IP routes
+ Dedicated circuits
= Significant performance gains



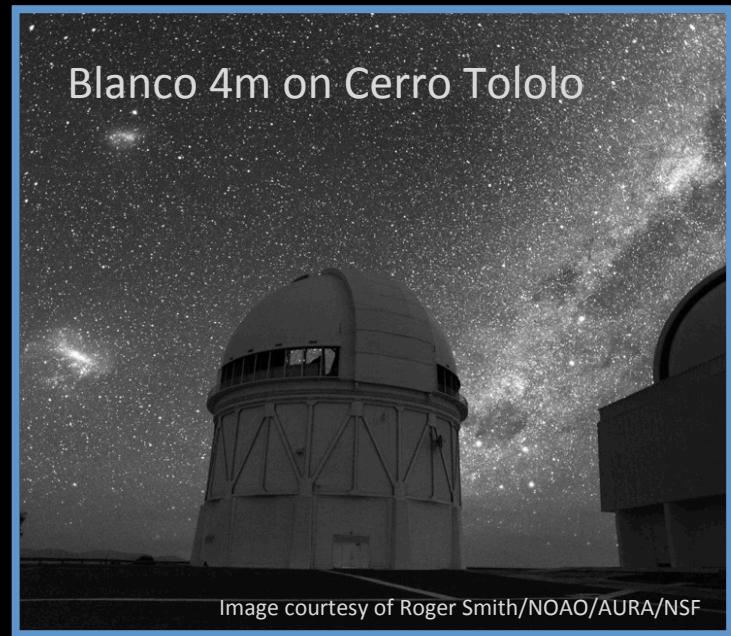
Provisioning OSCARS circuit for Globus Online transfers using Extensible Session Protocol (XSP)

Duration of runs, in seconds, over time.
Red: >10 TB transfer; green: >1 TB transfer.

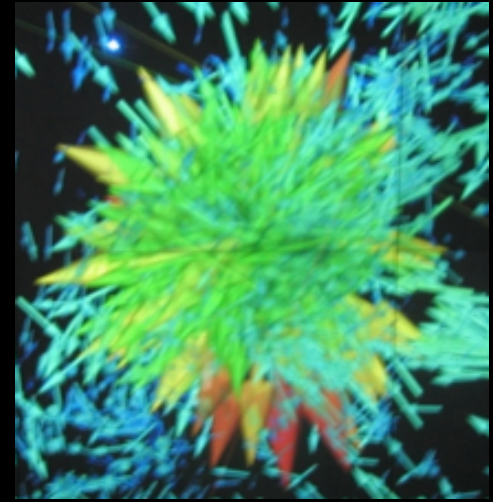
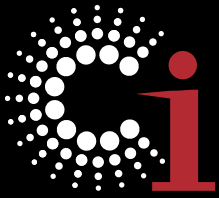




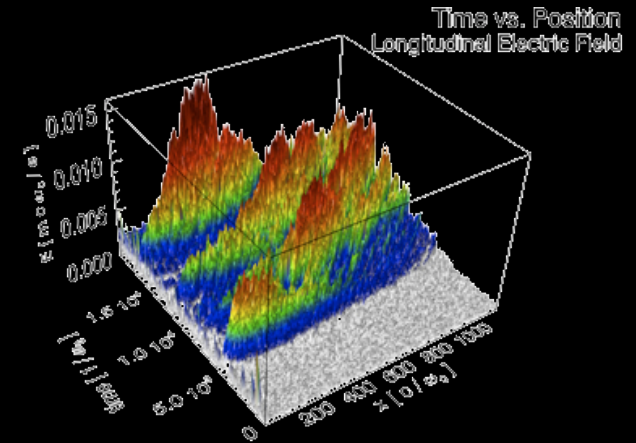
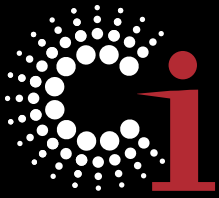
Dark Energy Survey



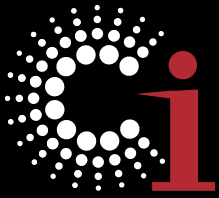
Receives 100,000 files each night in IL
Transmits files to Texas for analysis...
...then moves results back to IL
**Globus Online makes this process
reliable, routine, and efficient**



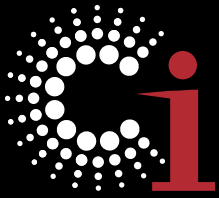
K. Heitmann (Argonne)
moves 22 TB of **cosmology**
data LANL → ANL at 5 Gb/s



B. Winjum (UCLA) moves
900K-file **plasma physics**
datasets UCLA → NERSC

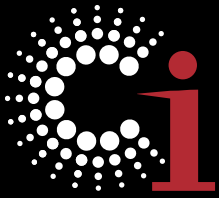


Dan Kozak (Caltech)
replicates 1 PB LIGO
astronomy data for resilience



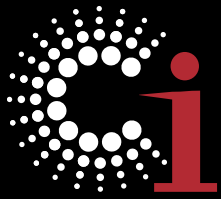
- Collect
- Move
- Sync
- Share
- Analyze
- Annotate
- Publish
- Search
- Backup
- Archive

...for **BIG DATA**



- Collect
- Move
- Sync
- Share
- Analyze
- Annotate
- Publish
- Search
- Backup
- Archive

...for **BIG DATA**



Many more capabilities planned ...

Globus Online Research Data Management-as-a-Service

**Ingest,
Cataloging,
Integration**

**Sharing,
Collaboration,
Annotation**

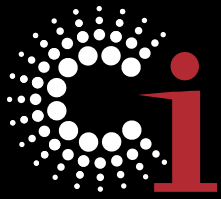
**Backup,
Archival,
Retrieval**

...

SaaS

Globus Integrate (Globus Nexus, Globus Connect)

PaaS



A platform for integration

KBase
PREDICTIVE BIOLOGY
DOE Systems Biology Knowledgebase

Home About News Developer Zone

The new Systems Biology Knowledgebase (KBase) is a collaborative effort designed to accelerate our understanding of microbial communities, and plants. It will be a community-driven, extensible source software framework and application system. KBase will provide access to data, models and simulations, enabling scientists to integrate new knowledge and share their findings.

Collaborate

What can KBase do?

- Combine heterogeneous data
- Offer standard interfaces
- Use evidence-based models
- Discover new insights
- Map genotypes to phenotypes
- Design and optimize
- Enable sharing

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

Reliable, high-performance, secure file transfer by Globus Online.

Blue Waters has partnered with the Globus Online file transfer service.

You may access this service by entering your Blue Waters username and password.

Sign In

Use Your NCSA Account

Earth System Grid

Home Data Account About Contact Us Logout

Globus Online Transfer: Step 2 of 3

Globus Online can be used to download the selected files to your local machine or to some other machine that has a GridFTP server. If you are downloading to local machine, you will need to do a one time setup of [Globus Connect](#), which can be downloaded from [Globus Online](#).

Please ensure Globus Connect is running before the next step.

Required Fields are Denoted by Blue text.

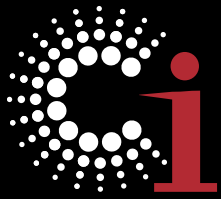
Destination Endpoint:

Destination Directory:

*nix: /tmp/
Windows: temp\

Next >> << Back Cancel

computationinstitute.org



Catalog as a service

Approach

- Hosted user-defined catalogs
- Based on tag model
 <subject, name, value>
- Optional schema constraints
- Integrated with other Globus services

Three REST APIs

/query/

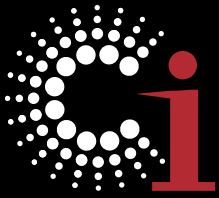
- Retrieve subjects

/tags/

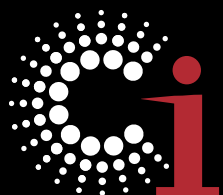
- Create, delete, retrieve tags

/tagdef/

- Create, delete, retrieve tag definitions



Other early successes in
services for science...



myexperiment^{beta}


myExperiment makes it really easy to find, use and share scientific workflows and other files, and to build communities.

All

Use myExperiment to...

- [Find Workflows](#)
- [Find Files](#)
- [Share Your Workflows and Files](#)
- [Create and Find Packs of Items](#)
- [Create and Join Groups](#)
- [Find People and Make Friends](#)
- [Send Messages](#)
- [Get Feedback](#)
- [Tag and Rate things](#)
- [Write Reviews and Comments](#)
- [Build your Profile and](#)

Explore



[About myExperiment](#)
[Join the Mailing List](#)
[Give us Feedback](#)
[For Developers](#)
[The myGrid Project](#)
[Taverna Workflow Workbench](#)
[The BioCatalogue Project](#)
[myExperiment Publications](#)

Register

or Login:

Username or Email:

Password:

Remember me:

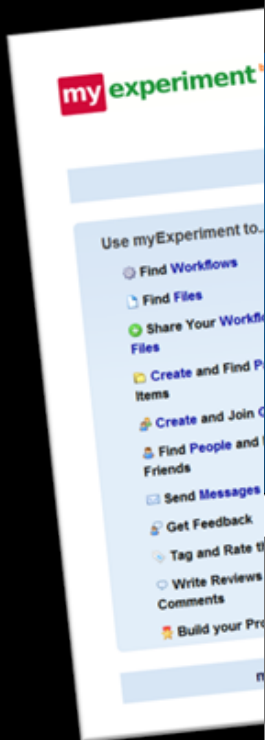
☐


Or use OpenID:

(eg. name.myopenid.com)

[Forgot Password?](#)

myExperiment has over 1250 users, 100 groups, 490 workflows, 130 files and 40 packs






an ncn project

ONLINE SIMULATION AND MORE
FOR NANOTECHNOLOGY

- Home
- My HUB
- Resources
- Members
- Explore
- About
- Support



FUNDAMENTALS OF NANOELECTRONICS

Online Course Spring 2012

[Learn More >](#)

◀ 1 2 3 ▶ || ▶


SIMULATE with over 160 tools for nanoelectronics, nanophotonics and more >

RESEARCH & COLLABORATE via groups, question board and more >

TEACH & LEARN with tool-powered curricula, courses, seminars and more >

SHARE & PUBLISH tools and research through our easy upload process

RESOURCES



Popular Tags: [nanoelectronics](#) [course lecture](#)

[material science](#) [Illinois](#) [nano/bio](#) [nanotransistors](#)

[research seminar](#) [devices](#) [nanophotonics](#)

[quantum transport](#) [tutorial](#) [transistors](#)

[molecular electronics](#) [nano electro-mechanical systems](#)


[NEGF](#) [carbon nanotubes](#) [nanomedicine](#)

[education/outreach](#) [UIUC](#) [band structure](#) [ABACUS](#)


[atomic force microscopy](#) [quantum dots](#) [MOSFET](#)

[nanowires](#) [More tags >](#)


★ FEATURED




[MIT Atomic Scale Modeling Toolkit](#) : Tools for Atomic Scale Modeling - in Tools



[Nanotechnology: Considerations for Facility Design](#) - in Online Presentations



[Greg Lush](#), University of Texas at El Paso - Contributions: 30



[Topics For Introductory Materials Classes](#) - in Topics



nanoHUB.org
an NCN project

ONLINE SIMULATION AND MORE
FOR NANOTECHNOLOGY



MG-RAST

metagenomics analysis server

LOGIN

REGISTER

PASSWORD

FORGOT?

login



Browse Metagenomes

search for metagenomes



Register



Contact



Help



Upload

*



News

About

MG-RAST (the Metagenomics RAST) server is an automated analysis platform for metagenomes providing quantitative insights into microbial populations based on sequence data.

# of metagenomes	35,586
# base pairs	9.24 Tbp
# of sequences	85.21 billion
# of public metagenomes	7,167

The server provides web based upload, quality control, automated annotation and analysis for samples up to 10GBp. Comparison between large numbers of samples is enabled via pre-computed abundance profiles. MG-RAST was launched in 2007 and has over 5000 registered users and 35,586 data sets. The current server version is 3.1.2.

Updates

MG-RAST Version 3.1.2 released

[Sign Up](#)[Login](#)[How It Works](#)

The Scientific Services Marketplace

The easiest way to get experiments conducted by researchers in top core facilities and institutions.

Immunohistochemistry **\$10.00**
per Sample

Cloning Constructs **\$162.00**
per Sample

NGS Sequencing **\$488.00**
per Sample

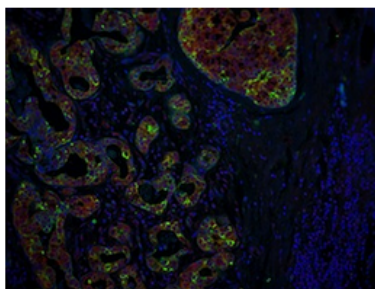
HTS Screening **\$0.10**
per Sample

Mass Spectrometry **\$10.00**
per Sample

RNA microarray **\$107.50**
per Sample

Featured Immunohistochemistry Providers

Science Exchange has 22 verified Immunohistochemistry providers including the following featured providers.



Histopathology and Tissue Shared Resource

Georgetown Lombardi Comprehensive Cancer Center | Washington, DC, United States

The Histopathology and Tissue Shared Resource at Georgetown Lombardi Comprehensive Cancer Center offers a full complement of histology services through combined Human Tissue Bank, histology, and IHC departments.

\$19.00 USD per Sample

[REQUEST ESTIMATE](#)

+17
others

[View →](#)

Free

No cost for providers to list services
• No cost for researchers to request services
• Easy to pay any provider via payment platform



Trusted

Trusted by researchers from leading institutions including Stanford, USC, Princeton and UC Davis • Great customer service by real people



Safe

Secure access to the world's best research infrastructure • Leading SSL and site encryption technologies to protect site information



From The Blog

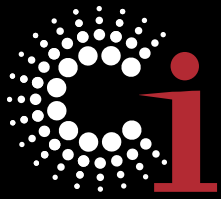
January 03, 2013

[Bringing Clarity to the Jargon](#)

[VISIT BLOG →](#)

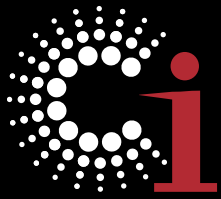
Updates

MG-RAST Version 3.1.2 released



Our vision for a 21st century cyberinfrastructure

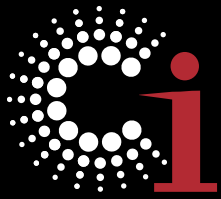
To provide **more** capability for **more** people at **substantially lower cost** by creatively **aggregating** (“cloud”) and **federating** (“grid”) resources —leveraging reliable, **high performance networks**



Our vision for a 21st century cyberinfrastructure

To provide **more** capability for
more people at **substantially**
lower cost by creatively
aggregating (“cloud”) and
federating (“grid”) resources

“Science as a service”



Thank you to our sponsors!



U.S. DEPARTMENT OF
ENERGY



THE UNIVERSITY OF
CHICAGO

Argonne
NATIONAL LABORATORY

